

Variability selected high-redshift quasars on SDSS Stripe 82

N. Palanque-Delabrouille¹, Ch. Yeche¹, A. D. Myers^{2,6}, P. Petitjean³, Nicholas P. Ross⁴, E. Sheldon⁵, E. Aubourg^{1,7}, T. Delubac¹, J.-M. Le Goff¹, I. Pâris³, J. Rich¹, K. S. Dawson⁹, D. P. Schneider¹⁰, and B. A. Weaver⁸

¹ CEA, Centre de Saclay, Irfu/SPP, F-91191 Gif-sur-Yvette, France

² Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana IL 61801, USA

³ Université Paris 6, Institut d'Astrophysique de Paris, CNRS UMR7095, 98bis Boulevard Arago, F-75014 Paris, France

⁴ Lawrence Berkeley National Lab, 1 Cyclotron Road, Berkeley, CA 94720, USA

⁵ Brookhaven National Laboratory, Bldg 510, Upton, NY 11973, USA

⁶ Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

⁷ APC, 10 rue Alice Domon et Léonie Duquet, F-75205 Paris Cedex 13, France

⁸ Center for Cosmology and Particle Physics, New York University, New York, NY 10003 USA

⁹ University of Utah, Dept. of Physics & Astronomy, 115 S 1400 E, Salt Lake City, UT 84112, USA

¹⁰ Department of Astronomy and Astrophysics, The Pennsylvania State University, 525 Davey Laboratory, University Park, PA 16802, USA

Received xx; accepted xx

ABSTRACT

The SDSS-III BOSS Quasar survey will attempt to observe $z > 2.15$ quasars at a density of at least 15 per square degree to yield the first measurement of the Baryon Acoustic Oscillations in the Ly- α forest. To help reaching this goal, we have developed a method to identify quasars based on their variability in the *ugriz* optical bands. The method has been applied to the selection of quasar targets in the SDSS region known as Stripe 82 (the Southern equatorial stripe), where numerous photometric observations are available over a 10-year baseline. This area was observed by BOSS during September and October 2010. Only 8% of the objects selected via variability are not quasars, while 90% of the previously identified high-redshift quasar population is recovered. The method allows for a significant increase in the $z > 2.15$ quasar density over previous strategies based on optical (*ugriz*) colors, achieving a density of 24.0 deg^{-2} on average down to $g \sim 22$ over the 220 deg^2 area of Stripe 82. We applied this method to simulated data from the Palomar Transient Factory and from Pan-STARRS, and showed that even with data that have sparser time sampling than what is available in Stripe 82, including variability in future quasar selection strategies would lead to increased target selection efficiency in the $z > 2.15$ redshift range. We also found that Broad Absorption Line quasars are preferentially present in a variability than in a color selection.

Key words. Quasars; variability

1. Introduction

Baryonic Acoustic Oscillations (BAO) and their imprint on the matter power spectrum were first observed in the distribution of galaxies (Cole et al., 2005; Eisenstein et al., 2005). They can also be studied by using the H I Lyman- α absorption signature of the matter density field along quasar lines of sight (White, 2003; McDonald & Eisenstein, 2007). A measurement sufficiently accurate to provide useful cosmological constraints requires the observation of at least 10^5 quasars, in the redshift range $2.2 < z < 3.5$, over at least 8000 deg^2 Eisenstein et al. (2011). This goal is one of the aims of the Baryon Oscillation Spectroscopic Survey (BOSS) project (Schlegel et al., 2009), part of the Sloan Digital Sky Survey-III¹ which is currently taking data. One of the challenges of this survey is to build a list of targets that contains a sufficient number of quasars in the required redshift range.

Quasars are traditionally selected photometrically, based on their colors in various bands (Schmidt & Green,

1983; Croom et al., 2001; Richards et al., 2004, 2009; Croom et al., 2009). While these methods achieve good completeness at low redshift ($z < 2$), they present serious drawbacks for the selection of quasars at redshifts above 2.2. In particular, as was shown in Fan (1999), quasars with $2.5 < z < 3.0$ tend to occupy the same region of optical color space as the much more numerous stellar population, causing the selection efficiency (or purity) to drop below $\sim 50\%$ in that region. The same confusion occurs again for $3.3 < z < 3.8$. This was recently confirmed by Worseck & Prochaska (2010) who have demonstrated that the SDSS standard quasar selection systematically misses quasars with redshifts in the range $3 < z < 3.5$.

The separation of stars and quasars in the redshift range of interest can be improved by using the variability of quasars in the optical bands. Light curves sampled every few days over several years were used by the MACHO collaboration (Geha et al., 2003) to identify 47 quasars behind the Magellanic Clouds. In a similar way, the OGLE project (Dobrzycki et al., 2003) has identified 5 quasars behind the Small Magellanic Cloud. Three seasons of obser-

¹ <http://www.sdss3.org>

vation on high galactic latitude fields were used by QUEST to search for variable sources. Nine previously unknown quasars (Rengstorf et al., 2004) were discovered.

More recently, significant progress in describing the evolution with time of quasar fluxes has been made possible by the multi-epoch data in the SDSS Stripe 82 (York et al., 2000). Using large samples of over 10,000 quasars, deVries et al. (2004) and MacLeod et al. (2008) have characterized quasar light curves with structure functions. Concentrating on SDSS Stripe 82 data, Schmidt et al. (2010) developed a technique for selecting quasars based on their variability. Recent works have shown that the optical variability of quasars could be related to a continuous time stochastic process driven by thermal fluctuations (Brandon et al., 2009) and modelled as a damped random walk (MacLeod et al., 2010a; Kozłowski et al., 2010). This resulted in a structure function that was used by MacLeod et al. (2010b) to separate quasars from other variable point sources. A variant, based on a statistical description of the variability in quasar light curves, was suggested by Butler & Bloom (2010) for the selection of quasars using time-series observations in a single passband.

In this paper, we present a method to select quasar candidates, inspired from the formalism developed by Schmidt et al. (2010). The method was adopted by the BOSS collaboration to choose the objects that were targeted, during September and October 2010, in Stripe 82. This region covers 220 deg^2 defined by equatorial coordinates $-43^\circ < \alpha_{J2000} < 45^\circ$ and $-1.25^\circ < \delta_{J2000} < 1.25^\circ$. It was previously imaged about once to three times a year from 2000 to 2005 (SDSS-I), then with an increased cadence of 10-20 times a year from 2005 to 2008 (SDSS-II) as part of the SDSS-II supernovae survey (Frieman et al., 2008). With a sampling of 53 epochs on average, over a time span of 5 to 10 years (Abazajian et al., 2009), the SDSS Stripe 82 data are ideal for testing a variability selection method for quasars. For the first time, in September and October 2010, the observational strategy of BOSS rested entirely on variability for the final selection (after loose initial color cuts as explained below). In contrast, all target lists in BOSS had been obtained so far from the location of the objects in color-color diagrams, following various strategies — such as the kernel density estimation method (Richards et al., 2004) or a neural network approach (Yeche et al., 2010).

Section 2 presents the formalism used to describe the variability in quasar light curves and gives the performance of the chosen selection algorithm on quasar and star samples. Section 3 explains how this tool was applied to select two sets of targets in Stripe 82, and presents the results obtained. An extrapolation of this method to the full $10,000 \text{ deg}^2$ observed by SDSS, made possible by adding data from the Palomar Transient Factory (Rau et al., 2009), or from Pan-STARRS ², is presented in Section 4. We conclude in Section 5.

2. Variability selection algorithm

The main purpose of this study was to develop an algorithm to select quasars in Stripe 82 based on their variability, while rejecting as many stars as possible. Spectroscopically confirmed stars and quasars in Stripe 82 were used to compute two sets of discriminating variables. The first one, used

to distinguish variable objects from non-variable stars, consists in the χ^2 of the light curve with respect to the mean flux, in each of the five photometric bands. The second one, which helps discriminating quasars from variable stars, consists in parameters that describe the structure function.

2.1. Quasar and star samples

We describe below the two samples, one of stars and one of quasars, which are used to test the variability algorithms, and to train the neural network of Sec. 2.5.

For the quasar training sample, we used a list of 13328 spectroscopically confirmed quasars obtained from the 2dF quasar catalog (2QZ; Croom et al., 2004), the 2dF-SDSS LRG and Quasar Survey (2SLAQ) (Croom et al., 2009), the SDSS-DR7 spectroscopic database (Abazajian et al., 2009), the SDSS-DR7 quasar catalog (Schneider et al., 2010) and the first year of BOSS observations. These quasars have redshifts in the range $0.05 \leq z \leq 5.0$ (cf. Fig 1) and g magnitudes in the range $18 \leq g \leq 23$ (Galactic extinction-corrected).




Fig. 1: Redshift distribution of the sample of quasars from all previous quasar surveys covering Stripe 82.

For the star sample, we used 2697 objects observed by BOSS, initially tagged as potential quasars from color selection and spectroscopically confirmed as stars. Variability and color-selection are not fully independent: bright objects that are easily discarded by their colors are also easier to discard by their variability. Therefore, the use of these spectroscopically confirmed stars constitutes a conservative approach and corresponds exactly to the type of objects that we want to reject with the variability algorithm.

Light curves were constructed for these two samples from the data collected by SDSS. The collaboration used the dedicated Sloan Foundation 2.5-m telescope (Gunn et al., 2006). A mosaic CCD camera (Gunn et al., 1998) imaged the sky in five *ugriz* bandpasses (Fukugita et al., 1996). The imaging data were processed through a series of pipelines (Stoughton et al., 2002) which performed astrometric calibration, photometric reduction and photometric calibration. Typical examples of stellar and quasar light curves are shown in Figs. 2 and 3

² <http://pan-starrs.ifa.hawaii.edu/public/home.html>




Fig. 2: Examples of light curves (after median filtering and clipping as explained in Sec. 2.2) in the five SDSS photometric bands for stars in Stripe 82.




Fig. 3: Examples of light curves (after median filtering and clipping as explained in Sec. 2.2) in the five SDSS photometric bands for quasars in Stripe 82.

respectively. The increased cadence after MJD 53500 are the SDSS-II supernovae search observations.

The star and quasar samples have similar time samplings, representative of the typical time sampling on Stripe 82 (cf. Figs. 2 and 3). The number of epochs (i.e. number of photometric measurements in a given band) varies from 1 to 140, with a mean of 53 and a r.m.s. of 20. The time lag between the first and the last epochs is 8 to 10 years long for 74% of the targets, between 5 and 7 years long for

24% and at most 4 years long for the remaining 2%. For this study, we concentrated on objects with at least 4 observation epochs, independently of the timespan. As a result, all targets that meet this requirement (13063 spectroscopically confirmed quasars and 2609 stars) have observations spanning at least two consecutive years.

2.2. Pre-treatment of the light curves

Photometric outliers could alter significantly the values of the variability parameters, to the point of washing out any relevant information. The raw light curves were therefore cleaned of deviant points (irrespective of their origin, whether technical or photometric) in a two-step procedure. A 3-point median filter was first applied to the full quasar light curve in each of the five bands, followed by a clipping of all points that still deviated significantly from a fifth order polynomial fitted to the light curve. Note that to avoid removing too many photometric epochs, the clipping threshold, initially set at 5σ , was iteratively increased until no more than 10% of the points were rejected. Despite the poorer frequency of the SDSS-I measurements (compared to SDSS-II), the median filtering was applied to the full light curve as the variations looked for are expected to occur on periods of several years.

2.3. Light curves χ^2

While most stars have constant flux, quasars usually exhibit flux variations. As shown by Sesar et. al. (2007), at least 90% of bright quasars are variable at the 0.03 mag level, and the variations in brightness are on the order of 10% on time scales of months to years (Vanden Berk et. al., 2004).

Each of the *ugriz* light curves were fit by a constant flux, and the resulting χ^2 recorded. While most stars have a reduced χ^2 near unity, as expected for non-variable objects, quasar light curves tend to be poorly fit by a constant, resulting in a large reduced χ^2 , as illustrated in Fig. 4 for the *r* band. The χ^2 thus helps to distinguish non-varying stars from varying point sources.

2.4. Variability structure function

The structure function characterizes light curve variability by quantifying the change in amplitude Δm_{ij} as a function of time lag Δt_{ij} between observations at epochs *i* and *j*. Following the prescription of Schmidt et al. (2010), the variability structure function of the source magnitude, is given by

$$\mathcal{V}(\Delta t_{ij}) = |\Delta m_{ij}| - \sqrt{\sigma_i^2 + \sigma_j^2}, \quad (1)$$

where σ is the magnitude measurement error. The structure function can be modeled by a power law $A(\Delta t)^\gamma$ in all photometric bands, with $\gamma > 0$, illustrating the fact that, for quasars, the r.m.s. of the distribution of the magnitude difference between two observations tends to increase with time lag (cf. Fig. 5).

To derive the power law parameters *A* and γ for a given light curve, we define the likelihood

$$\mathcal{L}(A, \gamma) = \prod_{j>i} \mathcal{L}_{ij}, \quad (2)$$




Fig. 4: Normalized distribution of the reduced χ^2 in the r band that results from fitting the light curves by a constant, for the stellar (blue) and the quasar (red) test samples. As confirmed by their larger reduced χ^2 , quasars clearly exhibit much larger deviations from a constant flux than stars.




Fig. 5: Variability structure function $\mathcal{V}(\Delta t)$ of equation 1 for a typical quasar. The curves show the best-fit power law $A(\Delta t)^\gamma$ for the three bands g , r , i . Note that the r and i best-fits are almost identical.

where for each ij pair of observations, an underlying Gaussian distribution of Δm values is assumed:

$$\mathcal{L}_{ij} = \frac{1}{\sqrt{2\pi\sigma^2(\Delta m)}} \exp\left(-\frac{\Delta m_{ij}^2}{2\sigma^2(\Delta m)}\right). \quad (3)$$

From the model above, the variability of the object, described by a power law, is naturally introduced in the definition of the variance $\sigma(\Delta m)^2$ of the underlying Gaussian distribution as

$$\sigma^2(\Delta m) = [A(\Delta t_{ij})^\gamma]^2 + (\sigma_i^2 + \sigma_j^2). \quad (4)$$

The A and γ parameters were then obtained by maximization of the likelihood $\mathcal{L}(A, \gamma)$ with the MINUIT package.³

We found that only the g , r and i bands had useful discriminating power: quasars have little flux in the u band

³ <http://wwwasdoc.web.cern.ch/wwwasdoc/minuit/min-main.html>

due to the Lyman continuum absorption of the intergalactic medium for rest frame wavelengths below 91.2 nm, and both u and z -band light curves exhibit more noise than the other light curves due to observational limitations (imaging depth and sky background variations in the u and z bands).

The fitted value of the γ parameter is roughly independent of the band. The fitted amplitudes in the different bands are strongly correlated but not identical. For instance, the g band amplitude is on average larger than the r band amplitude by about 0.04. To reduce the uncertainty on the fitted parameters, we therefore chose to fit simultaneously the g , r and i bands for a common γ and three amplitudes (A_g, A_r, A_i). We observe an excellent correlation between the amplitudes fitted with a common γ and those fitted with an independent γ per band, which implies that the data are indeed consistent with a unique power law valid for all bands.

The range of values obtained for stars and quasars are shown in Fig. 6. Non variable objects (mostly stars) lie near the origin of the graph, while quasars populate the region of larger A and γ values. It is interesting to notice that this approach can also distinguish various variable populations. RR-Lyrae, for instance, can have large variations (thus large A) but with no (or little) trend in time, implying that γ remains small. The necessary discrimination against variable stars, however, implies that quasars that exhibit a star-like variability cannot be found by this method. The same is even more true for non-variable quasars.




Fig. 6: Parameters γ and A_r of the variability structure function for the stellar (blue points) and quasar (red points) test samples. Large A 's indicate large fluctuation amplitudes. Large γ 's indicate an increase of the fluctuation amplitude with time.

2.5. Variability selection of quasars using a Neural Network

To complete our method for discriminating stars from quasars, an artificial Neural Network (NN) was used (Bishop, 1995).⁴ The basic building block of the NN architecture is a processing element called a neuron. The NN architecture used in this study is illustrated in Fig. 7,

⁴ We used a C++ package, TMultiLayerPerceptron, developed in the ROOT environment (Brun et al., 1995).

where each neuron is placed on one of four “layers”, with N_l neurons in layer l .




Fig. 7: Schematic representation of the artificial neural network used here with N_1 input variables, two hidden layers, and one output neuron.

The input of each neuron on the first (input) layer is one of the N_1 variables defining an object. Despite a lesser discriminating power of the u and z bands compared to g , r and i , the χ^2 's are robust quantities that can be used for all five bands. This is not the case for the structure function parameters, which result from a non-linear fit and were restricted to gri . Therefore, for the present study, the chosen variables are the four structure function parameters (γ , A_g , A_r and A_i) and the five χ^2 's, leading to $N_1 = 9$.

The inputs of neurons on subsequent layers ($l = 2, 3, 4$) are the N_{l-1} outputs (the x_j^{l-1} , $j = 1, \dots, N_{l-1}$) of the previous layer. The inputs of any neuron are linearly combined according to “weights” w_{ij}^l and “offsets” θ_j^l :

$$y_j^l = \sum_{i=1}^{N_l} w_{ij}^l x_i^{l-1} + \theta_j^l \quad l \geq 2. \quad (5)$$

The output of neuron j on layer l is then defined by the non-linear function

$$x_j^l = \frac{1}{1 + \exp(-y_j^l)} \quad 2 \leq l \leq 3. \quad (6)$$

The fourth layer has only one neuron giving an output $y_{\text{NN}} \equiv y_1^4$ reflecting the strength of quasar-like variability (as probed by the training sample) of the object defined by the N_1 input variables.

Certain aspects of the NN procedure, especially the number of layers and the number of nodes per layer, are somewhat arbitrary. They are chosen by experience and for simplicity. In contrast, the weights and offsets must be optimized so that the NN output, y_{NN} , correctly reflects the probability that an input object is a quasar. To determine the weights and offsets, the NN must therefore be “trained” with a set of objects that are spectroscopically known to be either quasars or stars. This is done with the test samples described in Sec. 2.1.

The result of the NN output is illustrated in Fig. 8. As expected, most stars peak near 0 while quasars usually have an output value near 1, and very few objects appear in




Fig. 8: Output of the variability Neural Network for the star and quasar samples. 97% of the quasars have $y_{\text{NN}} > 0.5$, and 3% are classified as star-like based on their variability ($y_{\text{NN}} < 0.5$). The histograms are normalized.

the middle range where the variability-based classification is uncertain.

Only 383 quasars out of 13063 (3%) are not classified as “quasar-like” by the variability NN, i.e. yield $y_{\text{NN}} < 0.5$. A visual inspection of their light curves confirms that they exhibit no clear variability, neither on short nor on long time-scales. A minimum loss of $\sim 3\%$ is therefore to be expected for any variability-based algorithm to select quasars using these data. This loss approaches 5% for the subsample of 3571 quasars at $z > 2.15$, probably due to the lower photometric precision of the objects. Part of the loss might also be due to the smaller rest frame time gap at high redshift.




Fig. 9: y_{NN} (color map) for the quasar sample, as a function of magnitude in g and number of epochs.

y_{NN} is independent of the time span. On average, for quasars, y_{NN} increases slightly with the number of epochs, as shown in Fig. 9, reaching its asymptotic value for about 40 epochs. It also depends on the object magnitude, with a shift of about 0.1 on average between $g \simeq 22.5$ and $g \simeq 18.5$. Most of the objects in Stripe 82 are well-sampled and

bright enough not to be affected by these small variations of performance. The results given hereafter are obtained after integration over the full distributions in magnitude and in number of epochs of the quasar sample.

To quantify the performance of our quasar selection, we define the completeness C and the purity P :

$$C = \frac{\text{Number of selected quasars}}{\text{Total number of confirmed quasars}}, \quad (7)$$

$$P = \frac{\text{Number of selected quasars}}{\text{Total number of selected objects}}. \quad (8)$$

We also define the stellar rejection R as

$$R = 1 - \frac{\text{Number of selected stars}}{\text{Total number of stars in the sample}}. \quad (9)$$

Fig. 10 illustrates the performance, in terms of quasar completeness and stellar rejection, of the variability-based NN, splitting the quasar sample in three different redshift ranges. For an identical stellar rejection, the loss of quasar completeness with increasing y_{NN} is enhanced at high redshift.




Fig. 10: Stellar rejection R vs. quasar completeness for the variability-based NN. Open circles are for known quasars at redshift z below 2.0, squares for those with $2.0 < z < 2.6$ and triangles for those with $z > 2.6$. Filled symbols at $R \sim 94.5\%$ and $R \sim 99\%$ indicate the location, on these curves, of the selection thresholds used in Sec. 3.1 and 3.2.

The small redshift-dependence of the variability-based selection method is further confirmed in Fig. 11, which shows the completeness $C(z)$ for the two thresholds on y_{NN} used in Sections 3.1 and 3.2. In contrast to a standard quasar selection based on colors, the completeness obtained here depends monotonously on redshift and has no minimum at any particular redshift. For a loose cut on the output of the variability NN ($y_{\text{NN}} > 0.50$), a high completeness is achieved at all redshifts. As the cut is tightened ($y_{\text{NN}} > 0.95$), however, a strong decrease with redshift appears, due to the reduced elapsed rest-frame time at high redshift, and to the decrease in the light curve signal-to-noise ratio as objects become fainter, resulting in a weaker significance of the variability. Nevertheless, even with a tight cut, the method still does not introduce any sharp redshift-specific feature.




Fig. 11: Completeness C vs. redshift for two thresholds on the output of the variability NN corresponding to those used for the selections of Sec. 3.1 (main sample, with $y_{\text{NN}} > 0.50$) and 3.2 (extreme variability sample, with $y_{\text{NN}} > 0.95$).

The purity of the selection cannot be determined as easily since it refers to a reference sample. The training sets are subsamples of the target population (they do not include, for instance, quasars selected through their variability but not through their colors). Knowledge of the total number of selected objects requires a complete sample of targets. Purity will therefore be given in Sec. 3.3, for two cases where the variability selection has been applied to actual data.

3. Variability-based selection on Stripe 82 for BOSS

BOSS is aiming at a density of $\sim 20 \text{ deg}^{-2}$ quasars at redshifts $z > 2.15$ (hereafter called “high- z ” quasars), with an allocation of 40 deg^{-2} optical fibers to obtain spectra of quasar candidates. In this context, the above study can be applied with two major goals.

The first one is to improve significantly the purity of the list of quasar candidates for which the spectra will be obtained. In BOSS, a traditional color-based selection with single epoch photometry typically reaches a quasar density of $10\text{--}15 \text{ deg}^{-2}$ from an initial selection of $\sim 40 \text{ deg}^{-2}$ targets. An algorithm with a higher purity presents the advantage of reaching the desired quasar density for BOSS, meaning an increase of about a factor 2, while keeping the number of fibers fixed. This is the aim of the “Main sample” described in Sec 3.1.

The second goal is to search for additional quasars, that would have been missed by previous searches because of colors beyond the typical range considered so far for quasars, but that could be selected based on their variability. This is the strategy leading to the selection of the “Extreme variability sample” presented in Sec. 3.2. These targets are expected to constitute a sample that would be less biased with redshift than through color selections. It would contribute to improving our knowledge of the quasar population in the approximate redshift range between 2 and 4.

Both approaches were adopted by BOSS for the observation of Stripe 82 in September and October 2010. The

results obtained are given in Sec. 3.3, and a comparison with color-based selections is presented in Sec. 3.4.

3.1. Main sample

The goal of the Main sample was to obtain a list of about 35 deg^{-2} targets with high quasar purity.

A color-based analysis with very loose thresholds is used to yield an initial list of $\sim 70 \text{ deg}^{-2}$ objects, expected to be dominated by stars by at least a ratio 2:1. Quasars are seen to have varying colors with time, since their structure function amplitudes A are band-dependent while the power γ is unique for all bands. However, the color change over a decade is observed to be small, with an average shift of 0.1 mag only. We thus co-added single epoch observations (cf. Fig. 12) to improve the photometry of the objects and their color measurements. The criteria for the preselection were defined as follows:

- output of a color-based NN > 0.2 (with colors determined from co-added observations) to remove objects that were far from the quasar locus in color-space (Yeche et al., 2010),
- $(u - g) > 0.15$ to enhance the fraction of $z > 2.15$ quasars over low- z ones. This cut rejects only 1% of previously known high- z quasars.




Fig. 12: Number of SDSS-I and SDSS-II measurements used to derive the co-added photometry in Stripe 82.

The completeness of this preselection for high- z quasars is of order 85%, which corresponds to an upper bound on the completeness of the “main sample”.

Requiring $y_{\text{NN}} > 0.50$, and removing previously identified low redshift quasars, we obtained a selection of 7586 objects (i.e. a target density of 34.5 deg^{-2}), called hereafter the “main sample”. Technical reasons related to the tiling of the objects (Blanton et al., 2003) reduced this sample to a density of 31.1 deg^{-2} . As shown in Fig. 11, the completeness of the variability selection at this threshold is expected to be $\sim 95\%$ (of the sample to which it is applied).

For comparison with the more usual color selection, we can remove the final variability selection and replace it by a tightened color cut (still using co-added photometry) adjusted to also produce a sample of 7586 targets. This color-selected sample and the main sample have 73% of their

targets in common. As clearly visible in Fig. 8, the threshold of 0.50 is very loose. There is thus no additional gain to be expected by lowering further the variability threshold.

Fig. 13 shows that the target density is flat with Right Ascension, as expected for extragalactic objects, in contrast to the peak that would be expected for $\alpha_{\text{J2000}} \simeq -43^\circ$ in the case of large contamination by Galactic stars as is seen in the initial distribution corresponding to the loose photometric preselection.

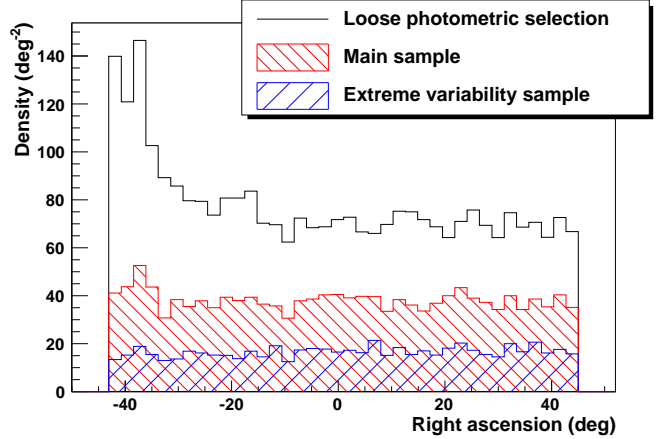


Fig. 13: Right Ascension distribution of targets in the main sample at the stage of loose color-based selection (black histogram), and after the final variability-based selection (red histogram). The targets of the extreme variability program are shown as the blue histogram.

Fig. 14 shows the distribution of the magnitude in the r band for the different samples. The drop at $r > 21$ is due to the color preselection. The selection leading to the main sample (red histogram) does not change the shape of the initial distribution (black histogram). This agrees with the fact that little redshift (and magnitude) dependence is observed at a threshold of 0.50 on the variability NN (cf. Fig. 11). The relative efficiency of the variability selection with respect to the preselected sample is roughly independent of magnitude.

3.2. Extreme variability sample

The second goal was to obtain an independent and complementary list of about 3 deg^{-2} objects selected by the variability NN but rejected according to their colors. With this approach, we could expect to find quasars in the stellar locus, at the risk of obtaining a sample dominated by variable stars rather than by quasars. This sample, however, offers a unique opportunity to explore a new region of color-space. Given the high level of discrimination between quasars and stars that is seen Figs. 6 and 8, the extreme variability sample is expected to have a strong potential.

The total number of point-like objects in Stripe 82 is on the order of several millions. Because the computation of the variability parameters on such a large sample would have been both disk- and time-consuming, a very loose preselection of about 1000 deg^{-2} objects was first applied, with the following criteria:




Fig. 14: Distribution of the magnitude in r at the stage of loose color-based preselection (black histogram), and after the final variability-based selection leading to the main sample (red histogram). The targets of the extreme variability program are shown as the blue histogram.

- $i > 18$ to limit the contribution from low- z quasars but $g < 22.3$ to maintain the possibility to obtain a good spectrum,
- $(g - i) < 2.2$ to exclude M stars,
- $(u - g) > 0.4$ to enhance the fraction of $z > 2.15$ quasars compared to low- z ones,
- $c_1 < 1.5$ or $c_3 < 0$ to remove a region in color-space distant from quasars and strongly populated by stars, where colors c_1 and c_3 are defined in Fan (1999) as

$$c_1 = 0.95(u - g) + 0.31(g - r) + 0.11(r - i),$$

$$c_3 = -0.39(u - g) + 0.79(g - r) + 0.47(r - i).$$

While these cuts reduced the total number of objects by about a factor of ten, leading to a sample of about 235,000 targets over the 220 deg^{-2} area of Stripe 82, they rejected only about 9% of previously known quasars at $z > 2.15$, uniformly over the magnitude range.

Requiring $y_{\text{NN}} > 0.95$ (i.e. selecting the most variable objects) then yielded a sample of 4360 targets (or a density of $\sim 20 \text{ deg}^{-2}$) called hereafter the “extreme variability sample”. Not all the targets could be observed: technical limitations (allocated number of fibers and tiling) reduced this sample to a density of $\sim 15 \text{ deg}^{-2}$.

The distribution of the Right Ascension of the selected objects is shown in Fig. 13 as the blue histogram. Its flatness is again an indication of low stellar contamination.

The magnitude distribution of this sample is illustrated in Figure 14 as the blue histogram: the selection efficiency drops by about a factor of two between the maximum, for a magnitude near 20, and its level at magnitudes near 22. This drop is to be expected given the decrease of completeness with redshift shown in Fig. 11.

About 65% of the extreme variability-selected quasars is also part of the main sample of Sec. 3.1. Because of the technical limitations mentioned above, which are tighter for the extreme variability sample than for the main one, the overlap increases to 78% of the actual targets. The remaining targets constitute what we call hereafter the “extreme variability *only* sample”. It contains 748 objects (i.e. a density of 3.4 deg^{-2}) for which spectra were measured.

3.3. Results

Thanks to good weather conditions, all planned targets have been observed. The reduction of the spectra was performed by the BOSS pipeline (Bolton & Schlegel, 2009), which also gives a preliminary determination of the redshift of the identified quasars. All spectra were then checked visually to yield final identifications and redshifts. Special features such as Broad Absorption Line (BAL) quasars were identified during this visual inspection. The pipeline and visual scanning are in agreement for $\sim 95\%$ of the objects. The spectra will be made available with the SDSS data release DR9, expected for mid-2012. A small selection is given in Fig. 15.

The outcome of the targeting of the two samples described above is summarized in Table 1. A total of 5270 high-redshift quasars were confirmed (4900 in the main sample, 2650 in the extreme variability sample of which 370 not in common with the main sample), a significant improvement over previous results. About half of these quasars (2770) were not known previously and were revealed by the present study. As stated in the abstract, we see that 90% of the known high-redshift quasar population is recovered by its variability, and that 92% of the selected targets are quasars (i.e., only 8% non-quasars). This high purity is in agreement with the flat Right Ascension distributions of the two samples shown in Fig. 13, indicating negligible stellar contamination.

The main sample has a quasar purity of 93% on average and 72% at a redshift $z > 2.15$. From this sample alone, the average density of $z > 2.15$ quasars over Stripe 82 has been increased from $\sim 15 \text{ deg}^{-2}$ from previous BOSS observations to 22.3 deg^{-2} .

It is remarkable that 86% of the objects in the “Extreme var. only” category, all rejected according to their colors, are quasars. Half of these, furthermore, are at $z > 2.15$. These results confirm the expected potential of the extreme variability program.

Considering the full sample selected from its extreme variability (i.e. including the candidates in the main sample that fulfilled the requirement $y_{\text{NN}} > 0.95$, cf. line “Extreme var.” of Table 1), we achieve an even higher purity: 96% of the objects are quasars, and 80% are at a redshift above 2.15. These results imply that variability is indeed an efficient tool for selecting quasars against all other variable sources.

The results for high-redshift quasars are also given split into two redshift bins. The drop of completeness with redshift expected from Fig. 11 for the extreme variability sample appears clearly. This sample, much more than the main sample does, selects preferentially quasars in the $2.15 < z < 3.0$ than in the $z > 3.0$ bin: the respective purities in the two bins are 68% and 11% for the extreme variability sample vs. 58% and 14% for the main sample.

The low fiber budget allocated to the Extreme variability program does not make the study of its completeness a relevant issue. However, we note that with a target density of only 3 deg^{-2} , the extreme variability program raised the high- z completeness of the main sample by $\sim 6\%$.

Fig. 16 shows the redshift distribution of the quasar samples selected through variability. As expected from the cut on $u - g$, most are at $z > 2.15$, corresponding to the requirements of BOSS. Fig. 17 shows that the additional quasars selected via extreme variability tend to preferen-




Fig. 15: Selection of quasar spectra from the variability targets, here shown smoothed over 9 Å. Upper and lower left: low- z quasars. Upper and lower middle: high- z quasars. Upper right: Broad Absorption Line high- z quasar. Lower right: high- z quasar displaying a Damped Lyman- α absorption.

Selection sample	Target density	All quasars		$z > 2.15$			$2.15 < z < 3.0$			$z > 3.0$		
		Density	$P(\%)$	Dens.	$P(\%)$	$C(\%)$	Dens.	$P(\%)$	$C(\%)$	Dens.	$P(\%)$	$C(\%)$
Main sample	31.1	29.0	93	22.3	72	84	18.1	58	86	4.2	14	76
Extreme var.	15.1	14.6	96	12.1	80	45	10.4	68	49	1.7	11	31
Extreme var. only	3.4	2.9	86	1.7	49	6	1.4	41	7	0.3	8	5
Total	34.5	31.9	92	24.0	69	90	19.5	56	92	4.5	13	81

Table 1: Density, purity P and completeness C of variability-based selections of quasar candidates. Densities are in deg^{-2} over an area of 220 deg^{-2} . Purity is the ratio of the density of the quasars in a given sample to the target density. Completeness includes all identified high-redshift quasars, whether from their color, variability, radio emission, etc. Column “Target” is for all candidates, “All quasar” refers to confirmed quasars independently of their redshift. Line “Extreme var.” includes both the extreme variability sample and the main sample targets that fulfilled the requirement $y_{NN} > 0.95$. Line “Extreme var. only” refers to objects rejected from the main sample due to their colors.

tially lie in the $2.5 < z < 3.0$ redshift range where color-based selections are known to be incomplete. This indicates that a pure variability-based selection can indeed contribute to the recovery of quasars lost during the color selection. The low number of quasars at $z > 3.4$ prevents firm conclusions from being drawn on this higher redshift range.

The location of the additional quasars in color-color space is presented in Fig. 18. There is no indication that they form a new class of quasars; instead, they appear to extend the quasar locus into the stellar locus in all color-color diagrams, as expected from synthetic models of quasar evolution (Fan, 1999). The completeness of the extreme-variability sample is quite low (cf. Table 1), so we can expect many more quasars than found here to be located in disfavored regions of color-space. High- z quasars are therefore probably even less well separated from the stellar locus than previously thought.

The fraction of Broad Absorption Line (BAL) quasars among the $z > 2.15$ quasars is seen to be higher in the sample selected for its extreme variability than in the main sample that includes stricter color cuts. Comparing the two non-overlapping “main” and “extreme var only” samples,

we have

$$\frac{\text{Number of high } z \text{ BAL quasars}}{\text{Number of high } z \text{ quasars}} = \begin{aligned} &7.0\% \pm 0.4\% \quad (\text{Main sample}) \\ &14.6\% \pm 1.8\% \quad (\text{Extreme var. only}) \end{aligned}$$

This seems to indicate that quasars affected by BAL features tend to fall outside the color regions that are generally favored by quasars.

3.4. Comparison with color selection

We compare the results obtained from this work to color selections of quasars. Two cases are studied below. The first one is a traditional color selection using single-epoch photometry. The large number of observations in Stripe 82, however, also permits a second approach using photometry obtained on co-added images, i.e. deeper frames and with a higher signal-to-noise, as was used for the color preselection of the main sample. A color selection on co-added images is expected to be much more complete than one based on single epoch observations.




Fig. 18: Color-color plots indicating the stellar (blue) and quasar (red) loci, as well as the position of the 370 high-redshift additional quasars rejected from their colors but selected through the variability neural network (extreme variability sample described in Sec. 3.2).

In both cases, we derived lists of 34.5 deg^{-2} targets as for the total variability-based selection (main and extreme variability samples) presented in this paper. We compared the outcome of these color-based selections to that of the variability-based one, using the full set of quasars identified on Stripe 82 from their color, variability or radio emission. The outcomes of the different selections are in the ratio 0.5:0.7:1 for the single-epoch color selection, co-added color selection and variability (this work) selection respectively. Fig. 19 shows the redshift distribution of the quasars recovered from the different samples. The dip around for $2.5 < z < 3.2$ in both color selections is clearly visible.

The advantage of variability might have been larger still with a greater ratio of the 34.5 deg^{-2} fibers allocated to the extreme-variability sample, since the latter has a higher purity than the main sample (cf. Table 1). As variability and colors seem to yield complementary samples (some quasars can be selected one way and not in the other), the most promising method would be to use both pieces of information simultaneously.

4. Use of external data and application to the full SDSS sky

Given the success of the variability-based selection in Stripe 82, it would be interesting to apply it over a much wider area in the sky. One possibility would be to use jointly data from SDSS (one or two photometric measurements over $10,000 \text{ deg}^2$) and forthcoming data from the Palomar Transient Factory (PTF) or Pan-STARRS 1 (PS1), which cover the same $10,000 \text{ deg}^2$ at several occasions over 3 to 5 years. A strategy based on these various data sets can be useful to future surveys like BigBOSS⁵ or LSST (LSST, 2009; Ivezić et al., 2008).

4.1. Extrapolation to PTF

Since December 2008, PTF has taken data in the R band at the cadence of one measurement every 5 nights (Rau et al., 2009). The images can be co-added to produce 4 deep frames per year of observation. Apart from Stripe 82, most of the area covered by SDSS was observed only once. The data available for quasar searches at the end of the PTF sur-

⁵ <http://bigboss.lbl.gov>




Fig. 16: Stacked redshift distribution of the confirmed quasars, where the histograms represent the number of quasars in each of the non-overlapping samples. The total extreme-variability sample is thus illustrated by the blue+purple surface, while the total main sample is in purple+red. The emphasis of the selection on $z > 2.15$ objects is apparent.




Fig. 17: Redshift distribution of the fraction of quasars added by the extreme variability selection compared to quasars in the same variability range but fulfilling color constraints.

vey can therefore be expected to consist typically of 1 point from SDSS (useful to extend the lever arm in time lag) and 4 points per year from PTF. To explore the possibilities offered by this data combination for quasar selection, we constructed synthetic light curves by down-sampling data from Stripe 82 in the following way:

- The last 5 years of SDSS are used to simulate PTF measurements: four evenly spaced points per year are selected from the SDSS data,
- To simulate the sole measurement available from SDSS on most of the sky, one point is taken at random over the previous years of SDSS, maintaining a gap of at least 2 years between the SDSS point and the first PTF measurement (to ensure a realistic lever arm).

Only synthetic light curves with all 21 measurements (1 for SDSS and 4 for each of the 5 years of PTF) are considered hereafter. With this constraint, we are left with 2248 (83%) stellar and 11456 (86%) quasar light curves (out of the initial samples described in section 2.1).

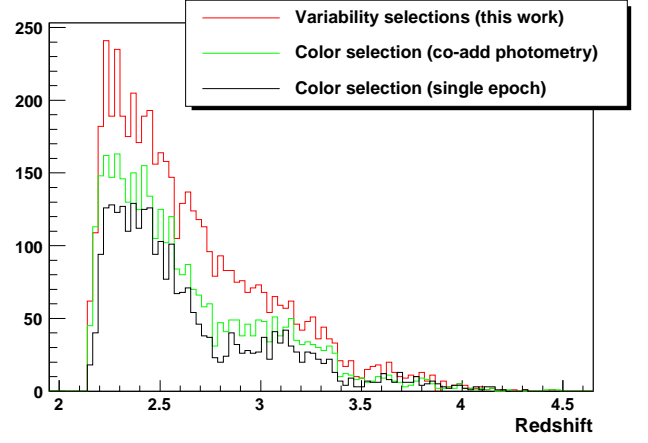


Fig. 19: Redshift distribution of the quasars recovered for three different selection algorithms presented in the text.

As PTF observes only in one band, the variability parameters are reduced to the reduced χ^2 in r , A_r and γ . A neural network was trained on the usual stellar and quasar test samples to yield an estimator of quasar likelihood based on these 3 parameters. The red triangles in Fig. 20 mark the evolution of the stellar rejection vs. quasar completeness as the threshold on the NN output is varied. They show that one can reach a quasar completeness of 85% for a rejection of 91% of the stars. For comparison, the blue dots illustrate the favorable case of Stripe 82 with all available measurements on 5 bands (case studied in Section 3) and a variability selection based on the 9-parameter NN.

Note that as explained in Sec. 2.1, the stellar sample used for figure 20 has passed loose color cuts that might not be available for PTF data. We have checked that the performance of the algorithm in the rejection of randomly picked Stripe 82 objects, statistically dominated by stars, is within 1% of the performance plotted in the figure.




Fig. 20: Stellar rejection vs. quasar completeness for the full Stripe 82 data (blue dots), for the Pan-STARRS (green and black squares) and for the PTF (red triangles) simulated data. In each case, the threshold on the relevant variability NN is increased from right to left.

4.2. Extrapolation to PS1

Pan-STARRS 1 (PS1) started regular observations in March 2009. With its 3 degree field of view, the whole available sky is recorded 3 times during the dark time of each lunar cycle. The first part of the project is expected to last about 3 years, after which a second telescope will begin operation. To explore the use of the PS1 data, we proceeded in a similar way as for PTF. The main difference is that PS1 has data available in five filters (g , r , i , z and y) instead of one. For quasar selection in the redshift range $2.15 < z < 4$, we considered only the filters in common with SDSS (g through z). This restriction produced 8 variability parameters: four χ^2 's (one in each of the four bands), A_g , A_r , A_i and the common γ (as for the study of Stripe 82). As for PTF, a NN was trained to yield an estimator based on these 8 parameters. The performance of the resulting selection is illustrated in Fig. 20 for two survey durations, 3 or 5 years. Only synthetic light curves with all 13 (in the case of a 3-year survey) or 21 (in the case of a 5-year survey) measurements are considered in the plot.

The 3-year survey gives results comparable to those for the 5-year PTF. In contrast, the 5-year PS1 survey is a significant improvement over the 3-year survey, and can reach an 85% quasar completeness for a 97% stellar rejection, or a 91% quasar completeness for a 95% stellar rejection.

The absence of the SDSS anchor point would reduce the quasar completeness by about 3%. Of course, the SDSS data would have little impact on the stellar rejection R , since most stars exhibit flat light curves, whatever their coverage.

4.3. Extrapolation to fainter high- z targets with PS1

Quasar selection was typically concentrated at $g < 22.3$. Future surveys like BigBOSS intend to go deeper in order to increase the density of quasars. To study the impact of a deeper magnitude limit on the performance of the variability selection, we used all objects defined as point sources in coadded frames to compute stellar rejection vs. quasar completeness for magnitude limits $g < 21$, $g < 22$ and $g < 23$, in the case of five years of PS1 data. The coadded images are used to detect the sources out to $g > 23$, while the lightcurves are still simulated by downsampling the shallower, single-epoch, SDSS data. The redshift range of interest for ground-based Lyman- α BAO studies is restricted to $z > 2.15$. In this section, we concentrate on these high- z quasars.

To extrapolate to fainter targets, the stellar sample is now taken to be a set of random objects in a 7.5 deg^2 region in Stripe 82 around $\alpha_{J2000} = 0$. It contains about 1000 objects per deg^2 at $g < 21$, and ~ 2500 at $g < 23$. The quasar sample is the one used before augmented by the new quasars discovered in Stripe 82 using the work presented in this paper (Sec. 3.3). We use it to compute the efficiency of quasar recovery in three non-overlapping magnitude bins: $g < 21$ (about 11000 quasars), $21 < g < 22$ (over 5000 quasars) and $22 < g < 23$ (about 2000 quasars). This sample is highly incomplete for faint objects. Therefore, to compute results integrated up to a given magnitude limit, we weight the efficiencies in each magnitude bin by a theoretical quasar luminosity function (LF) based on Hopkins et al. (2007) and extrapolated to low luminosities (cf. LSST science book). We also use the quasar LF (corrected by de-

tection efficiencies) to estimate the quasar contamination in the so-called stellar sample. This contamination is negligible in the original sample dominated by stars, but as the threshold on y_{NN} increases, actual quasars contained in the “stellar” sample begin to dominate the set of selected objects. To compute the rejection levels, their contribution is thus estimated and removed. We estimate the systematic uncertainty on the stellar rejection due to this correction to be of order 1%.

Fig. 21 shows the stellar rejection R as a function of quasar completeness C for high- z quasars. At 80% quasar completeness (respectively 90%), the stellar rejection decreases by $\sim 3\%$ (resp. 8%) when changing the limit from $g < 21$ to $g < 23$.




Fig. 21: Stellar rejection vs. $z > 2.15$ quasar completeness for five years of Pan-STARRS simulated data. The depth varies from $g < 21$ (blue dots) to $g < 23$ (red triangles). The green dot-dashed lines show the effect of a factor of 2 uncertainty on the quasar LF for $22 < g < 23$. The red dotted lines illustrate the effect of a $\pm 5\%$ change on the quasar recovery efficiency. The red dashed line indicates the expected improvement by combining variability and color selections. From left to right, points correspond to $y_{NN} = 0.95, 0.92, 0.87, 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, 0.2, 0.1$ and 0 .

To study the impact of the uncertainty on the quasar LF, we varied the LF in the $22 < g < 23$ magnitude bin by a factor of 2 either way. As shown in Fig. 21 (green dot-dashed lines), even such a large change of the LF has little impact on the results (1% at most). This is because the quasar recovery efficiency decreases only moderately with increasing magnitude, in particular at large quasar completeness (for y_{NN} near 0.6 or less).

The uncertainty on the $g < 23$ curve is dominated by the uncertainty on the recovery efficiency for quasars in the $22 < g < 23$ magnitude bin. In this study, it is estimated at a mean magnitude $g \sim 22.3$, lower than what is expected from the quasar LF. The impact of a 5% uncertainty on this efficiency is shown in the figure as the red dotted curves: after integration over the full magnitude range, a 5% change in the recovery efficiency shifts the $g < 23$ curve by about 2%.

The stellar rejection for a given completeness of high-redshift quasars can be improved significantly by combin-

ing variability and photometric criteria. With a similar approach as what was done for BOSS on Stripe 82, we define main and extreme variability samples using photometric information from BOSS single-epoch data. The only photometric cut for the main sample is $\mathcal{P}_{\text{ED}} > 10^{-3}$, where \mathcal{P}_{ED} is the probability of extreme deconvolution defined in Bovy et al. (2011). This cut rejects 4% of the high- z known quasars. About half of these can be recovered with the extreme variability sample, defined by $y_{\text{NN}} > 0.95$ and loose photometric cuts similar to those applied on Stripe 82 (Sec. 3.2). The resulting performance is shown in Fig. 21 as the upper red dashed line (for all objects up to $g < 23$). Considering the $g < 23$ curve, relevant to future surveys, we obtain a stellar rejection $R = 99\%$ for a quasar completeness $C = 80\%$, and $R = 98\%$ for $C = 90\%$. Variability alone would have yielded instead $R = 95\%$ and $R = 90\%$ respectively in the same $z > 2.15$ redshift range. In addition, the photometric selection is optimized for the rejection of low- z quasars, whereas variability is not.

Although the variability method cannot lead to results as good for the sparser data of Pan-STARRS (13 to 21 measurements in four bands) or PTF (21 measurements in one band) as for the SDSS data on Stripe 82 (~ 50 measurements in five bands), it can still contribute significantly to quasar selection. Used in addition to a color selection, as was done with BOSS for Stripe 82, even with a single epoch in SDSS (for areas other than Stripe 82), it results in much improved selections than what color-selection alone can achieve.

5. Conclusions

We have designed a method that characterizes light curve variability in order to discriminate quasars from both non-variable and variable stars. A Neural Network was implemented to yield an estimator of quasar likelihood derived from these variability parameters.

The method has been applied in conjunction with a loose color-based preselection to define a list of 31 deg $^{-2}$ targets in Stripe 82 for which spectra were taken with BOSS. The performance of this selection on quasars at redshift above 2.15 can be quantified by a purity of 72% and a completeness of 84%. This represents a significant improvement over traditional fully color-based selections which seldom obtained a purity in excess of 40%.

A second study was dedicated to the objects exhibiting an extreme quasar-like variability. An additional 3 deg $^{-2}$ targets were selected on the following criteria: the objects had to be excluded from the previous sample (i.e. did not have favorable colors according to quasar standards), and had a very high value of the output of the variability NN. Half of the selected objects proved to be high redshift quasars and 40% low redshift quasars. This program thus increased further the completeness of the quasar selection, reaching the unprecedented value of 90% total on average over Stripe 82.

Combining the above two programs allowed BOSS to obtain a density of $z > 2.15$ quasars in Stripe 82, all selected through their variability, of 24.0 deg $^{-2}$, with only ~ 35 deg $^{-2}$ fibers dedicated to their identification.

The method developed here was also applied to ersatz data from Palomar Transient Factory or from Pan-STARRS to determine the performance that can be achieved for fu-

ture target selections of quasars over about 10,000 deg $^{-2}$ of the sky.

Acknowledgements. Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy. The SDSS-III web site is <http://www.sdss3.org/>.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, University of Cambridge, University of Florida, the French Participation Group, the German Participation Group, the Instituto de Astrofísica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, New Mexico State University, New York University, the Ohio State University, the Penn State University, University of Portsmouth, Princeton University, University of Tokyo, the University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

ES is supported by grant DE-AC02-98CH10886. The BOSS French Participation Group is supported by Agence Nationale de la Recherche under grant ANR-08-BLAN-0222.

References

- Abazajian K. et al., 2009, ApJS, 182, 543
- Adelman-McCarthy J.K. et al., 2008, ApJS, 175, 297
- Bishop, C. M., “Neural Networks for pattern recognition”, 1995, Oxford University Press
- Blanton, M.R., et al., 2003, AJ, 125, 2276
- Bolton, A.S. and Schlegel, D.J., 2010, PASP, 122, 248
+ BOSS pipeline
- Bovy, J. et al., 2011, ApJ, 729, 141
- Brandon, K.C., Bechtold, J. and Siemiginowska, A., AJ, 2009, 698, 895
- Butler, N. and Bloom, J., 2010, arXiv:1008.3143
- Brun, R. et al. (the ROOT Team), <http://root.cern.ch>
- Cole, S. et al., (the 2dFGRS Team), 2005, MNRAS, 362, 505
- Croom, S.M. et al., 2001, MNRAS, 322, 29
- Croom, S. M. et al., 2004, MNRAS, 349, 1397
- Croom, S. M. et al., 2009, MNRAS, 392, 19
- de Vries, A., et al., 2004, AJ, 129, 615
- Dobrzycki, A., Macri, L., Stanek, K. and Groot, P., 2003, AJ, 125, 1330
- Eisenstein D. J. et al., (the SDSS Collaboration), 2005, ApJ, 633, 560
- Eisenstein D. J. et al., (the SDSS Collaboration), 2011, submitted to ApJ, arXiv:1101.1529
- Fan, X., 1999, AJ, 117, 2528
- Frieman, J. et al., 2008, AJ, 135, 338
- Fukugita, M., Ichikawa, T., Gunn, J.E., Doi, M., Shimasaku, K., and Schneider, D.P. 1996, AJ, 111, 1748
- Geha, M. et al., 2003, AJ, 125, 1
- Gunn, J.E., et al. 1998, AJ, 116, 3040
- Gunn, J.E., et al. 2006, AJ, 131, 2332
- Hopkins, P.F., Richards, G.T. & Hernquist, L., 2007, ApJ, 654, 731
- Ivezic Z., Tyson J. A., Allsman, R. et al. (LSST Collaboration), arXiv:0805.2366
- Kozłowski, S. et al., AJ, 2010, 708, 927
- LSST Science coll., 2009, version 2.0, arXiv:0912.0201
- MacLeod, C.L. et al., 2008, arXiv:0810.5159
- MacLeod, C.L. et al., 2010a, ApJ, 721, 1014
- MacLeod, C.L. et al., 2010b, arXiv:1009.2081
- McDonald, P. & Eisenstein, D., 2007, Phys. Rev. D., 76, 063009
- Nugent, P., private communication
- Percival, W. et al., 2010, MNRAS, 401, 2148
- Rau, A. et al., 2009, PASP, 121, 886
- Rengstorf, A. et al., 2004, ApJ, 606, 741
- Richards, G.T. et al., 2004, AJS, 155, 257
- Richards, G.T. et al., 2006, AJ, 131, 2766
- Richards, G.T. et al., 2009, AJS, 180, 67
- Schneider, D.P. et al., 2010, AJ, 139, 2360
- Schlegel, D., White, M., and Eisenstein, D., 2009, arXiv:0902.4680
- Schmidt, M. and Green, R.F., 1983, ApJ, 269, 352
- Schmidt, K.B. et al., 2010, ApJ, 714, 1194
- Sesar, B. et al., 2007, AJ, 134, 2236
- Stoughton C. et al., 2002, AJ, 123, 485

Vanden Berk, D. et al., 2004, ApJ, 601, 692

White, M., 2003, arXiv:0305474

Worseck G. and Prochaska J.X., 2010, arXiv:1004.3347

Yeche, Ch. et al., 2010, A&A, 523, 14

York, D.G., et al. 2000, AJ, 120, 1579